

SOME REMARKS ABOUT CONNECTIONISM IN PSYCHOLOGICAL SIMULATION

Alberto Greco

Department of Philosophy and Laboratory of Experimental Psychology
University of Genova, Italy

Abstract

This paper considers the question of whether the connectionist approach is an appropriate paradigm for the construction of simulation models of psychological processes. The strongest connectionist claim is that psychological properties simply emerge from interactions among the units of a network; meaningful features are distributed in these units, which therefore are not representations of anything. The problem is then: how such a system can describe a psychological system and be relevant to explain it ?

It is suggested that the psychological relevance of connectionist descriptions should not be rejected on the grounds that they are limited to a hardware-like or implementational level. Connectionist descriptions would be best considered as different scientific "objects", constructed from the same pre-scientific "facts". As different objects they need not be at different levels of generality or abstraction and the translation of one representation into another is unnecessary.

As to the explanatory power of connectionist models, it is claimed that it cannot come from the interaction among units, if this interaction is invoked at the same time to give the model a psychological nature. Neither can the explanatory power come from the fact that connectionist models produce phenomena similar to human psychological processes, since this similarity should refer explicitly to features that have already been accepted independently as psychological.

The claim that connectionist analysis is psychologically relevant at a sub-symbolic level is eventually considered. There is a possible confusion in this claim: sub-symbolic models may be non-symbolic as to their object but they must be symbolic as models. A non-symbolic model simply would not be a model, since - by definition - it must represent something. Sub-symbolic models, however, may be non-symbolic in the sense that there are no symbols in the modelled reality. In this regard, connectionist systems seem to help in overcoming the idea - encouraged by current AI and Cognitive Science - that any psychological process should be considered as made out of symbols.

Revised version of a paper presented at the 7th Workshop of the ESSCS, St. Maximin-la-Sainte-Baume, Provence, France, 19-22 June 1989.

1. Introduction

The purpose of the present paper is to examine whether the connectionist approach is a good paradigm in order to construct simulative models of psychological processes.

The connectionist paradigm has been discussed up to now specially as regards the modelling of cognitive processes, mainly because so far a solid system of simulations is available only of these processes. In this paper we prefer to focus on "psychological" phenomena, in order to consider the impact of connectionism on a broader scope, that is not limited to the use of knowledge but open to the problem of the genesis of knowledge. In addressing this question, it will become clear that the connectionist approach seems to support particularly the importance of taking into account non-cognitive aspects (i.e., not limited to conscious knowledge) to explain cognition, an already emerging issue in cognitive science (e.g. see Abelson, 1979; Norman, 1980; Miller, 1981; Schoenfeld, 1983).

Some general questions concerning the use of simulation in psychology will be examined first; in the following sections the explanatory power and the psychological relevance of some typical connectionist claims will be discussed.

2. Psychological Simulation and Psychological Objects

Psychological simulation is a methodology consisting of the construction of models of psychological events, which enable us to reproduce some aspects of them, using a programmed computer or whatever other system. The simulation method in psychology has the same goal as the experimental method, that is to understand (describe, explain) psychological phenomena.

Generally speaking, a psychological simulation should refer to a psychological object. The simplest definition of a psychological object is: what is accepted as such among psychologists. This definition may appear unsatisfactory at first sight, and in a sense it certainly is, but it has some merit other than its simplicity: it stresses the importance of defining what is the object of any discipline.

To better understand some points which are made below, it is worthwhile to briefly specify what kind of epistemology of the object construction has been adopted in this paper.

Any scientific discipline starts from some "fact" which is originally defined in terms of pre-scientific common sense. Different disciplines describe the same fact in different ways, from different points of view, and using different predicates, operational definitions, instruments of measure, etc. In short, they construct different objects. When describing some "fact" - concerning human mental activities, or human behaviors, or other phenomena recognized as being of interest by psychologists - a psychological object is then constructed

if a psychological viewpoint is adopted, that is concepts or predicates developed and accepted inside the psychological field are used.

3. Problems of Connectionist Models

In the classical view, a psychological simulation is conceived as a system in which there is a data structure processed by a computer program. The data are to be interpreted as symbols referring to psychological constructs ("concepts", "schemata", "images" and so on); the program is assumed to be isomorphic with some analogous natural way of processing these data, which also can be interpreted as some psychological construct (we can label it as a "mental process" or give it more specialized names such as memory, learning, vision, etc.). For example, in a typical problem solving simulation system there is a direct correspondence between the symbols processed by the program and the elements which constitute the problem from the psychological point of view (goals, operators, states, etc.).

Connectionist models, instead, do not consider programs and data but only networks of interacting units which can exhibit psychological properties such as learning, conceptualization, etc. These models aim to be simulations of psychological processes, although they are not always explicitly presented as such.

Since there are different positions inside current connectionist approaches (see Smolensky, 1988; Frixione, Gaglio, and Spinelli, 1989) it is appropriate to specify that here we are referring to the most typical kind of connectionist models, where units in a network do not represent anything and meaningful features are distributed in it. The term "strong connectionism" will be used to refer to this view.

The strongest connectionist view claims that psychological properties can simply emerge from interactions among the elements of a network. The "raw material" used are only strengths of connections among elements, and a labeling of these as excitatory or inhibitory. That is almost all. In particular, the nature of elements themselves does not matter. Since in this view the elements need not be viewed as symbols, they cannot correspond to psychological primitives, or anyway have a psychological nature. Even when connectionists are forced to claim that these elements correspond to "features" relevant to the phenomenon under simulation, they do not see the need to give these features the flavour of identifiable, real psychological objects. Thus the trouble with strongest connectionist models is that, although they are concerned with psychological reality, their elements are not symbols referring to any psychological concept as was the case in traditional psychological modelling. The main problem is whether or how is it possible to model a psychological system without using symbols that have a psychological character (or without using symbols at all).

4. The explanatory power of connectionist models

The first question to ask is: why should a system of non-specific interacting elements explain a psychological or cognitive system? What does it add to our understanding of how this system works, of which principles are behind it ?

A possible answer to this question could be: because it works, it produces a phenomenon analogous to some human process. Connectionists do not state it directly this way, in such naive terms, but undoubtedly they often underline the psychological plausibility of their models. For example, vis-a-vis one connectionist device, one is tempted to say that it sees something because it behaves as humans do when they see something, perhaps it avoids some obstacle, it grasps the right thing, and so on. Suppose (in the best case, of a partially local network) that we have constructed this device such that it can put together some features corresponding to graphemic, orientational, spatial elements: does this simple fact enable us to take as proven that this is the way humans see as well?

Apart from the philosophical question of "intentionality", this reason could be sufficient, but not necessary. If it were a necessary reason, there would be little trouble in believing, for example, that any automatic retrieval system is a model of human representation or recognition: to recall an old saying, even a key entering the keyhole would be considered as "recognizing". This is a crucial point. A machine that retrieves is not a machine that represents: therefore, the connectionist use of terms such as "representational features" or "knowledge atoms" might be inappropriate.

If it is true, as emerges from these considerations, that doing the job rightly is not enough, what more is needed exactly?

5. The Analogy Requirement of Models

In answering the question whether a model is or is not adequate in giving explanations in some scientific discipline, one cannot talk about how a model explains something without knowing what it is a model of. In other terms, one cannot rely on what this model does but should look at what its object is. A more general discussion, again, is needed on this subject.

It is important to point out that simulating a phenomenon means to give rise to a new phenomenon. This could be the utterance of different words, a set of numbers, a system of graphs, the operation of physical devices, etc. This new phenomenon - which we call a "model" - is, of course, not equal to the original but similar to it in some respects. One can understand the real phenomenon insofar as this correspondence works as a hypothesis, a theory about it and since at every moment it is possible to shift from the model to what it is supposed to represent.

This method, however, is of use only if the aspects which have been singled out really play a crucial part in the original phenomenon. The problem for

any simulation is how to guarantee this correspondence and the commonest solution is to point out some analogy between the model and the modelled phenomenon. This requirement is important since a model includes some features essential to the model's working but not to the understanding of the modelled object (this is the so-called negative analogy, Hesse, 1966).

It is clear, therefore, that in constructing a model of some phenomenon, similarities or analogies between the model and its object must be exactly specified. But it should be also clear that for a psychological model this specification should be accomplished, again, adopting a psychological viewpoint. For example, one feature of a model can be described in many ways, as a number or as an electric current or, say, as an operator in a problem solving space; evidently, not all of these descriptions can be approached from a psychological viewpoint.

6. Explanation and Construction of the Object

We have seen above that different descriptions involve different objects. This is important also for the explanation aspect, since different descriptions involve different explanations or, in some cases, no explanations at all. To give a trivial example, if you describe an electrical installation by analyzing all its wires from a chemical point of view, you cannot explain why, on pushing a switch, the lamp lights up. You have to adopt the relevant point of view. Likewise, (as we have seen) a hardware description of what happens inside a computer when we are running, say, an expert system program cannot be used to explain why the system fails to solve some problem, and so on. The same happens with a favourite connectionist example: in Newtonian and quantum mechanics there are simply different objects.

The question we were asking above (why should a system of non-specific interacting elements explain a psychological or cognitive system), is then how to recognize some psychological object, in the sense mentioned above, in these interacting elements. For example, what enables us to define, say, as a "concept" the result of the interaction among some assemblies of elements where nothing is a concept (or some other psychological object)? If it is only a label that has been attached to the product of the system, recognized a posteriori as working as a concept, then it is difficult to see how this could be a truly psychological object (note 1).

So the question remains: where does the psychological nature in connectionist systems come from? In answering this question, the problem is that if you analyze strong connectionist models, on one side, at the units level, you cannot find anything corresponding to psychological units; but on the other side, if you claim that their psychological nature arises from the interaction among units, and, going back to your memories of Gestalt, you say that the whole is more than the sum of its parts, then you run into another problem: your

object acquires a psychological nature only after a process, which happens to be the same process that is invoked to explain the real phenomenon.

For example: what makes, say, "vision" what some model does? Where can visual patterns be found in it? The connectionist answer is that they cannot be found in single units, but are distributed; these patterns emerge as a result of interaction. Why then do we see the way we do? The answer is: because this interaction takes place in the way showed by the model.

In other words, the interaction among elements is what enables you to construct your object as psychological and, at the same time, what enables you to explain it. Which does not seem completely correct. To explain should mean to give "reasons why" and these reasons should be independent from reasons alleged in constructing the object.

7. The Question of Levels and Nature of Connectionist Objects

The problem of the nature of connectionist descriptions has been apparent to the connectionists themselves from the beginning. In 1985, in their paper on distributed memory (subsequently inserted in the PDP volume), McClelland & Rumelhart pointed out that their approach, compared with traditional ones, aimed to provide both a different description and a different level of analysis of cognitive phenomena. In his well-known comment, Broadbent (1985) argued that only the latter was the case, namely that connectionist models are relevant only at the low, hardware-like level, the one called implementational by Marr (1982), and that they had nothing to say at a more general level of description, the computational one.

This has become the most typical criticism of the psychological relevance of connectionist models. But it seems quite odd that connectionists themselves usually accept a comparison of their approach with others simply in terms of "levels". Even in the latest connectionist formulation, such as Smolensky's "Proper Treatment of Connectionism" (1988), this comparison is still made in terms of "levels of analysis".

The idea that connectionist descriptions are relevant at a different level in comparison with traditional ones, however, could be a wrong presupposition. Where a different level of analysis is concerned, probably a different description (more or less general, abstract, close to physical, etc.) of the same phenomenon is meant. However, connectionists insist that their models are alternative to, and not simply implementations of, traditional models.

We can better understand this problem in which connectionists themselves are often involved if we consider again how scientific objects are constructed from pre-scientific common sense. We have seen that a unique "fact" can be described using different disciplinary instruments. These different descriptions, however, are not necessarily at different levels; if they are assumed to be, then a further assumption must be made: namely, that there is a hierarchy of representations, or that these descriptions are hierarchically arranged from

top to bottom or vice versa (for example the bottom may be less general, less abstract, closer to neurophysiological reality, and so on). But this assumption is unnecessary and perhaps misleading.

In fact, there are some advantages in considering different descriptions as different points of view instead of different levels. The main one is that you need not bother with the problem of translation from one level to another. Consider the classical "mind-body problem", which well illustrates this point: you can describe the same (pre-scientific) event either at a neurophysiological or psychological level, but in fact if you try to translate one of them into the other you get stuck. This is because the concepts that you use when speaking of the brain have no equivalent in the "mind" domain. The computer metaphor has been useful, at least, to understand that a description of what happens when running a computer program, made in terms of electric currents in the circuits, is not directly translatable into a description in terms of flow-charts or algorithms of the very same event, occurring in the very same time and place. Thus, as we have stated above, the different point of view simply constructs another thing; in fact we are dealing with different objects. We could speak of peculiar "connectionist objects".

So it would not make any sense to ask the question whether psychological interest lies more in high-level, intermediate, or low-level descriptions (such as Marr's "computational", "representational" and "implementational" levels). Again, a description is psychological and you can speak of psychological phenomena only from a psychological viewpoint, namely if the description makes use of psychological primitives or concepts accepted in a psychological context.

To sum up: the connectionist claim that distributed "representations" have psychological relevance at an intermediate level between the mental level and the neural level would be best replaced by the claim of having produced, from the peculiar "connectionist viewpoint", new objects. These objects correspond to one description of some artifacts, made using connectionist concepts (i.e. as networks of units with strength values of connections, working according to some equations, etc.). At this stage, extending this description and saying that these artifacts are able to recall, infer, etc. (notably, without making use of symbols) is possible exactly as it is possible, metaphorically, to say that a chip in a computer "remembers" how to solve a problem. Two problems, indeed, remain to go beyond this metaphorical attribution of model properties to such artifacts: whether the use of symbols is an essential property of any model and what could make the connectionist viewpoint have some relevance in explaining psychological phenomena. We will discuss these problems in the next two sections.

8. The Symbolic Nature of Connectionist Models

To better analyze what correspondence there might be between connectionist models and psychological phenomena, let's return to the issue of what connectionist models represent (or whether they represent anything at all). We

have seen above that, in any scientific context, the elements that constitute a simulation model must represent some aspects of the real phenomenon; in other words, they are symbols (note 2) that stand "in place of" the original phenomenon (or in place of features of it) (note 3). For example, in cognitive psychology we may consider the elements of a simulation model as symbols representing psychic states (pieces of knowledge, concepts, representational structures like frames, etc.). In other areas of psychology, we could also think of these symbols as representing feelings, moods, fears, ideas, etc.

From the strongest connectionist position, it would seem that these statements no longer hold. For example, in a network that "categorizes" a concept there is nothing representing the concept, not even something that represents the categorization itself.

We have already seen that the claim of having produced some phenomenon resembling a psychological one does not necessarily mean that the original has been simulated. In any case, this implies specifying such similarities, and this cannot be done without making use of symbols.

In fact, even before constructing the model one already must have a symbolic description of the phenomenon to be simulated. Usually this description is provided by existing theories and it consists of linguistic statements about the phenomenon considered. According to the epistemology above outlined, this description is relevant to a discipline (namely to psychology) only if it adopts its peculiar viewpoint. After its construction the model, being a new phenomenon, can be symbolically described on its own. To do so, again, different disciplinary viewpoints can be used and it can be considered on its own, in every respect, a new object.

Since a model is not the real object, what relationship can it have with this object? One cannot see what other kind of relationship there could be if not of referring to it. In this sense a model cannot be non-symbolic.

However, connectionists typically say that they are making a psychologically relevant analysis of the real phenomenon, not at a symbolic level, but at a sub-symbolic level.

But what about the sub-symbolic consideration of models? Are they symbolic or not symbolic? Admittedly, their elements are not, since they do not represent anything. But can we say that the overall model "represents" some psychological process? Neither, because, as we have seen, the working system does not seem to stand for the simulated process but rather it seems to produce it, to make it! Thus we have a system that shows some ability to remember, to categorize, to recognize, etc. but we are not told why we should believe that the way they do it (that is, as a result of interactions among elements following some specified mathematical or statistical function) is the way people do it (note 4).

But here we must avoid a possible confusion: the fact that a simulation system must be symbolic since it must represent some real phenomenon does not mean that in principle the nature of that phenomenon should be symbolic as well. You may simulate phenomena in which you could not find any symbol,

even if to do so you must use some symbols. What in my opinion causes such confusion about sub-symbolic connectionist models is that they are symbolic as models but not necessarily symbolic as to their object.

Let's examine what this means. Firstly, what does it mean to say that connectionist models are symbolic as models? As we have seen, it is impossible to devise a model which is not symbolic in some respects: it simply would not be a model. Therefore, the connectionist claim that their models do not make use of symbols cannot be interpreted as the claim that absolutely nothing in these models is symbolic.

In fact, in building and using such models there is necessarily some point at which one must use symbols; at least, when the overall model is interpreted: the interpretation cannot be disjoined from the model itself and it must be present from the beginning of the construction of the model (note 5). When working with their networks and when talking about them, connectionists are obliged to read some psychologically meaningful and symbolic activities - such as "recognition", "learning", "comprehension" - into their models. To say that a model learns, remembers, behaves in a certain way, etc., one must have some criterion to assess it and to check if it does the things it is intended to do. This criterion must be symbolic, in the sense that the aspects shared by the model and by the real object must be described using some symbolic system. To claim that a phenomenon has been reproduced is not to claim that the real thing has been produced; in any case one must symbolically specify what aspects of the latter have been captured and where the correspondence between the two systems lies.

Secondly, connectionist models may be non-symbolic as to their object. This is a very different issue from the previous one: in this case, it can be said that there are no symbols in connectionist models since there are no symbols in the modelled reality. But a full understanding of this point involves careful consideration of the nature of a psychological object in simulative models.

9. The Symbolic Nature of Psychological Objects

We have seen that, in general, the object of a simulation (what phenomena can be simulated, what the model refers to) substantially coincides with whatever may be accepted as a legitimate object of a discipline. Is a symbolic nature a necessary requirement for any psychological object?

Even if it is not questioned that a symbolic nature is essential to some psychological processes, such as language, for some others this is less sure. Current artificial intelligence and cognitive science have perhaps overemphasized the symbolic aspect of psychological processes, mainly considering in which language they can be expressed.

One of the reasons the idea that psychological processes are best considered as symbolic systems has been encouraged is that cognitive psychology, during its growth in the seventies, has adopted (perhaps implicitly) the claim that

internal phenomena are operations, or rather processing activities, and the pervasive concept of "information" has quickly spread. So the idea of Newell & Simon of intelligence as a physical symbol system popularized the conception that psychological processes are themselves examples of symbol manipulation. Therefore the development of artificial intelligence and cognitive science has been grounded on a view of mental processes as rational activities (logical, linguistic, systematic, conscious). In fact, the simulation objects of inner states have almost always been conceived of as a sequence of states of conscious experiences: knowledge, concepts, rules, frames or scripts (or even beliefs, feelings, fears and so on, in some simulations as in Colby's or Abelson's early works). In this conception the syntactic and semantic aspects have been strongly separated and the relationship between a symbol and its meaning has been conceived as being rigorously arbitrary.

Internal psychic states, however, can be conceived legitimately as non-symbolic events: that is, not as referring to some experience, standing in place of it, but being the experience itself. For example, you could rightly suppose that when a person sees something there is nothing in his mind representing what he is seeing; seeing (or, more probably, some aspect of it) might well be conceived as a direct experience, possible without the use of symbols (this is, for example, the approach of Gibson, 1979, to visual perception). This is even truer, of course, in the case of feelings or insights. You do not have to think necessarily of mysterious or deep feelings, but, for example, only of the commonest problems of intermediate steps or of "insight" in problem-solving.

Now, if one wishes to reproduce some non-symbolic aspect of a psychic activity, there is an important difference between the traditional and the connectionist simulative approaches.

One of the greatest limits of traditional simulation, made by running programs on serial computers, was the computability requirement. Von Neumann computers can execute any task, provided that it can be described in a non-ambiguous language, stating precisely which rules to follow to go from one system state to the next. In principle, however, this does not mean that only activities whose nature is symbolic can be simulated, but that only activities for which symbols can be found to express them, are simulable. The essential requirement for the construction of any model is the possibility of finding symbols to express the phenomenon, not that the phenomenon itself should be considered symbolic.

In the first part of this paper I have tried to maintain that the interest of connectionist models is not in the fact that they are able to perform psychological functions or to work as psychological systems. It should be clear from that discussion that this is not a sufficient condition in order to explain these systems. In particular, connectionists fail to point out where the analogies between the two systems lie.

By contrast, there is another aspect of connectionist models which makes them interesting from a psychological point of view. This is the claim of showing how symbolic events emerge from non-symbolic events. I think there is no

reason why this claim should not be taken seriously. The connectionist perspective, in contrast to traditional cognitivist approach, makes it clear that symbols need not be considered as ready-made products. This perspective promises, at least, to cast light on the process of their construction, the "symbol formation" (after Werner & Kaplan, 1963).

In this perspective it is possible to construct truly psychological models, not only because it is in the historical tradition of psychological research but mainly because these models meet the analogy requirement. A connectionist system which discovers concepts, for example, may not be suitable to explain the process of concept discovery, because the analogy between the two processes, as a whole, cannot be found. However, this system could be suitable to show how a relation of representing/represented between some events is established, analogously to what happens in the human process of representation.

This topic, however, needs a wider discussion. We should be aware, for example, of the fact that accepting this perspective entails having a particular conception of human representation and symbolization (i.e., representation as a process which develops in time and symbolization as not entirely arbitrary). In addition, other problems must be resolved, such as whether or how these systems can have a self-organizing character. A full discussion of these questions is not possible here, it will be done elsewhere. Here it has been sufficient to outline the possibility of adopting the connectionist viewpoint to achieve models of the emergence of symbolic activities.

10. Conclusions

The question whether connectionism is a suitable paradigm in order to construct simulative models of psychological phenomena has been examined.

I have claimed that a psychological simulation model should refer to a psychological object, that is to a description of some pre-scientific "fact" made adopting a "psychological viewpoint", i.e. using predicates or concepts considered as acceptable among psychologists.

Differently from traditional psychological models, in connectionist models there are networks with interacting elements which do not have any particular meaning nor can be viewed as symbols referring to any psychological phenomenon. The problem is then whether this interaction can be considered a model of psychologically relevant phenomena.

A possible solution would be that the psychological interest of these models lies in the fact that they perform operations analogous with ones performed by humans (e.g. they learn). It must be considered however, that the usefulness of model construction lies in the specification of which analogies there are between one model and its object and that this specification also should be made adopting a "psychological viewpoint". In fact, the claim that connectionist systems explain how or why humans exhibit the same performance (e.g.,

learning) cannot be grounded if they fail to show the similarities between model and reality.

Connectionist models usually also explain psychological phenomena invoking the interaction among their internal units. However it does not seem completely correct to claim, at the same time, that this interaction gives the model a psychological nature.

The most popular solution to the problem whether connectionist systems are adequate descriptions of psychological processes has been that they are relevant at a low level, close to the neural one. Connectionists themselves seem to accept this statement. However, the consideration of psychological models at different levels could be misleading because it leaves open the problem of how to translate from one level to another and of the correspondence between low-level models of some process and high-level models of the same process. This is because when thinking of levels as hierarchically arranged it is natural to think of lower levels as implementations of higher ones.

I have suggested, as a solution to these problems, that connectionist models, instead of being considered the description at a different level of what happens in a system exhibiting psychological properties, would be best considered as descriptions of a different object, constructed from a peculiar "connectionist viewpoint". There are, then, two problems: how psychological modelling could be done without making use of symbols, and what could give this viewpoint psychological relevance.

Since a model is not the real object, it can be constructed and used only in a symbolic fashion, namely specifying for any of its features which of them refers to which psychological one. In this sense a model cannot be non-symbolic. In any case connectionists, proposing their models as psychological, are forced to use psychologically "pregnant" symbols (such as "stimulus", "recognition", etc.) in the overall model interpretation.

However, the claim of necessarily having symbols in models does not mean that the object of simulation must be symbolic. Cognitive science seems to have encouraged the idea that psychological processes consist only of symbolic activities, but connectionist systems could help in overcoming this idea. In particular, they seem suitable to simulate the process of symbol genesis.

Notes

1. A similar problem occurred with some early simulations of personality processes, such as Colby's simulation of paranoia (Colby, 1981). In such models, the idea (perhaps ingenuous) was to conceive mental phenomena as variables (in the programming language sense) with a label. For example, Colby's simulation aimed to show that paranoia is connected with humiliation. In this model, humiliation was simply a variable with a numerical value, which was increased or decreased according to what happened to other variables. But what

made it possible to call this variable "humiliation", instead of, say, "fear"? Only the label that Colby, relying on his own intuition, had attached to it.

2. In the present discussion the terms "symbol" and "symbolic" are used to mean any physical event to which it is possible to ascribe some meaning, that is which can be viewed as standing in place of something different from itself. Thus I have stressed the representational aspect of the symbolization which, I think, is the only aspect about which everyone can agree. I have not considered other properties of symbolic systems, such as being compositional or manipulable, which are often discussed with reference to connectionist models, since, in my opinion, they go beyond the original and simplest meaning of the concept.

3. This means using a system of symbols to represent it. Of course, you may use natural language, or other symbolic systems like formalized languages (mathematical, logical, statistical, etc.) or - in the case of simulation - artifacts arranged in such a way that you can recognize that what happens in, say, a pipeline system is similar in some respect to what happens to an economical system or to a social system, and so on. Even in this case the artifact system acts as a symbolic system.

4. These problems, of course, are not so difficult with "less strong" kinds of connectionism. For example, in distributed symbolic models, network units represent "micro-features": in this case, in the construction of the network one is forced to select some features that are considered "important" to characterize the phenomenon under simulation. These features usually do have a meaning. As an example, to simulate learning of the past tense of English verbs, Rumelhart & Mc Clelland (1986) had to invent the "Wickelphones" and the "Wickelfeatures" to represent the most important features of linguistic forms and to allow discriminability, generalizations, etc.

5. On this point the objection that the symbolic aspect does not reside in the model but in the model builder's intentions has been raised, and that the interpretation of models is not part of the models themselves. Here I believe there is a misunderstanding on what a model is: to model originally means to give a shape to something shapeless, following a design which is present before the model construction. It is difficult to see how could it be possible having a "model" without knowing what it is a model of.

References

- Abelson, R.P. (1979), Differences between belief and knowledge systems, *Cognitive Science* 3, 355-366.
- Broadbent, D.A. (1985), A question of levels: Comment on McClelland and Rumelhart, *J. of Experimental Psychology: General* 114, 189- 192.
- Colby, K.M. (1981), Modeling a paranoid mind, *The Behavioral and Brain Sciences* 4, 515-560.
- Frixione, M., Gaglio, S., & Spinelli, G. (1989), Symbols and subsymbols for representing knowledge: A catalogue raison\$, 11th IJCAI, Detroit, Mich., USA.
- Gibson, J.J. (1979), *The ecological approach to visual perception*, Boston: Houghton.
- Hesse, M.B. (1966), *Models and analogies in science*, Notre Dame, Indiana: University of Notre Dame Press.
- Marr, D. (1982), *Vision*, San Francisco: Freeman.
- McClelland, J.L., Rumelhart, D.E., & The PDP Research Group (1986), *Parallel distributed processing: Explorations in the microstructure of cognition*, Cambridge, MA: Bradford Books/MIT Press.
- McClelland, J.L., & Rumelhart, D.E. (1985), Distributed memory and the representation of general and specific information, *J. of Experimental Psychology: General* 114, 159-188.
- Miller, G.A. (1981), Trends and debates in cognitive psychology, *Cognition* 10, 215-225.
- Norman, D.A. (1980), Twelve issues for cognitive science, *Cognitive Science* 4, 1-32.
- Rumelhart, D.E., & McClelland, J.L. (1986), On learning the past tenses of English verbs, In J.L. McClelland, D.E. Rumelhart, & The PDP Research Group (1986), Volume 2.
- Schoenfeld, A.H. (1983), Beyond the purely cognitive: belief systems, social cognitions, and metacognitions as driving forces in intellectual performance, *Cognitive Science* 7, 329-363.
- Smolensky, P. (1988), On the proper treatment of connectionism, *The Behavioral and Brain Sciences* 11, 1-74.
- Werner, H., Kaplan B. (1963), *Symbol formation*, Wiley, New York.

Manuscript received: 26-9-1989, in revised form 19-3-1990.

Address author:

Dept. of Philosophy, Laboratory of Experimental Psychology,
University of Genova, Via Balbi, 4, 16126 Genova, Italy.